arXiv:1405.2682v1 [physics.comp-ph] 12 May 2014

Calculating vibrational spectra with sum of product basis functions without storing full-dimensional vectors or matrices

Arnaud Leclerc[1, 2, a)] and Tucker Carrington[1, b)]
[1)] *Chemistry Department, Queen's University, Kingston, Ontario K7L 3N6, Canada*
[2)] *Université de Lorraine, UMR CNRS 7565 SRSMC, Théorie-Modélisation-Simulation, 1, boulevard Arago 57070 Metz, France*

We propose an iterative method for computing vibrational spectra that significantly reduces the memory cost of calculations. It uses a direct product primitive basis, but does not require storing vectors with as many components as there are product basis functions. Wavefunctions are represented in a basis each of whose functions is a sum of products (SOP) and the factorizable structure of the Hamiltonian is exploited. If the factors of the SOP basis functions are properly chosen, wavefunctions are linear combinations of a small number of SOP basis functions. The SOP basis functions are generated using a shifted block power method. The factors are refined with a rank reduction algorithm to cap the number of terms in a SOP basis function. The ideas are tested on a 20-D model Hamiltonian and a realistic CH_3CN (12 dimensional) potential. For the 20-D problem, to use a standard direct product iterative approach one would need to store vectors with about 10^{20} components and would hence require about 8×10^{11} GB. With the approach of this paper only 1 GB of memory is necessary. Results for CH_3CN agree well with those of a previous calculation on the same potential.

I. INTRODUCTION

The most general and systematic way of solving the time-independent Schrödinger equation to compute vibrational bound states, and hence a vibrational spectrum, requires computing eigenvalues of a basis representation of the Hamiltonian operator. When standard methods of "direct" linear algebra are used to diagonalize the Hamiltonian matrix the memory cost of the calculation scales as N^2, where N is the size of the matrix and the number of basis functions. Diagonalization can be avoided by using an iterative eigensolver to compute the eigenvalues of interest. The Lanczos[1,2] and Filter Diagonalization methods[3–6] are popular iterative options. Iterative approaches require only the evaluation of matrix-vector products. If it is not possible to do matrix-vector products without keeping the Hamiltonian matrix in memory then the memory cost of iterative methods also scales as N^2. Fortunately, one can often exploit either structure of the basis (and the Hamiltonian operator) or sparsity of the Hamiltonian matrix to evaluate matrix-vector products without storing (and sometimes without computing elements of) the Hamiltonian matrix.[7–9] In both cases, the memory cost of a product basis iterative calculation scales as $N = n^D$, where n is a representative number of basis functions for a single coordinate and D is the number of dimensions, which is the size of a vector (at least two vectors must be retained in memory). Exploiting the structure of a product basis also makes it possible to evaluate matrix-vector products efficiently (at a cost that scales, regardless of the complexity of the potential, as n^{D+1}).[8–18]

The combination of iterative algorithms and product basis sets reduces the memory cost to n^D, which is obviously much less than n^{2D}, required to store the Hamiltonian matrix in the product basis. However, even n^D is very large if $D > 6$. If $D = 12$ and $n = 10$ then for a single vector one needs $\sim$ 8000 GB of memory. This is a manifestation of the "curse of dimensionality". To further reduce the memory cost one option is to use a better basis. It is not hard to devise a better basis, what *is* tricky is finding a basis that has enough structure to make it possible to efficiently evaluate matrix-vector products. If matrix-vector products are not evaluated efficiently, reducing the number of basis functions required to represent the wavefunctions of interest can (significantly) *increase* the CPU cost of the calculation. There are four popular ways of reducing basis size. First, prune a standard product basis by retaining only some of the functions.[19–27] Second, use contracted basis functions obtained by solving reduced-dimension eigenproblems.[28–35] Third, optimize 1D functions with the multi-configuration time dependent Hartree method (MCTDH).[36,37] Fourth, use basis functions localized in the classically allowed region of phase space.[38–41] Some pruned bases are compatible with efficient matrix-vector products, for both problems with simple potentials (requiring no quadrature)[42] and with general potentials (for which quadrature is necessary)[17,43–45]. Several ideas have been proposed for evaluating matrix-vector products with contracted bases.[46–49] Because MCTDH uses a direct product basis it is straightforward to evaluate matrix-vector products at a cost that scales as n^{D+1} (with n the number of single-particle functions). To date no one has attempted to use iterative methods in conjunction with phase space localized bases.

In this paper we propose an iterative method, for computing spectra, that significantly reduces the memory cost of calculations. We use a direct product basis (al-

[a)] Electronic mail: Arnaud.Leclerc@univ-lorraine.fr
[b)] Electronic mail: Tucker.Carrington@queensu.ca

though the ideas would also work with a pruned basis). To represent a wavefunction, all previous product-basis iterative methods store n^D coefficients. Our new approach is motivated by the realization that, in some cases, the n^D coefficients, used to represent a function, can be computed from a much smaller set of numbers. For example, a product of functions of a single variable, $\phi_1(q_1)\phi_2(q_2)\cdots\phi_D(q_D)$, can be represented as

$$\sum_{i_1=1}^{n} f_{i_1}^{(1)}\theta_{i_1}^{1}(q_1)\sum_{i_2=1}^{n} f_{i_2}^{(2)}\theta_{i_2}^{2}(q_2)\cdots\sum_{i_D=1}^{n} f_{i_D}^{(D)}\theta_{i_D}^{D}(q_D)$$

and it is only necessary to store Dn numbers. Obviously, for a real problem the wavefunction is not a product of functions of a single variable, but it should be possible to represent many wavefunctions as sums of products of functions of a single variable. If, for one wavefunction, R terms are required, one must store RDn numbers. This may be much less than n^D. When $n = 10$ and $D = 12$, $RDn < n^D$ if $R < 8\times10^9$. For many molecules it is surely possible to find a sum of products (SOP) representation of wavefunctions with a value of R small enough that it is worth exploiting the SOP structure to reduce the memory cost.

We develop a method using SOP basis functions to find eigenpairs of a SOP operator in section II. The memory cost scales as nRD, which is the memory required to store one SOP basis function, where R is the required number of terms in a SOP basis function. The key idea is to use basis functions that are sums of products of optimized factors. Basis functions are determined, from matrix-vector products evaluated by doing 1-D operations, by applying the Hamiltonian to other SOP functions. The number of terms in the basis functions is controlled by a reduction procedure. The reduction is a crucial part of the method we propose. In section III, the method is tested on multidimensional coupled oscillator models with D as large as 20. The lowest transitions of acetonitrile, CH_3CN, (a 12-D problem) are computed and compared with results of Avila et al[50] in section IV.

II. SUM OF PRODUCTS (SOP) EIGENSOLVER

A. SOP basis functions and CP format representation

Our goal is to calculate eigenstates of a Hamiltonian operator by representing it in an efficient SOP basis. We define a primitive product basis using 1-D functions $\theta_{i_j}^{j}(q_j)$ with $i_j = 1,\ldots,n_j$ for each coordinate q_j. The primitive basis is unusably large. An SOP basis function, $\Psi_k(q_1,\ldots,q_D)$, can be expanded in the primitive basis as

$$\Psi_k(q_1,\ldots,q_D) \simeq \sum_{i_1=1}^{n_1}\cdots\sum_{i_D=1}^{n_D} F_{i_1 i_2\ldots i_D}\prod_{j=1}^{D}\theta_{i_j}^{j}(q_j)\ . \quad (1)$$

For SOP basis functions,

$$F_{i_1 i_2\ldots i_D} = \sum_{\ell=1}^{R}\prod_{j=1}^{D} f_{i_j}^{(\ell,j)}\ , \quad (2)$$

where $f^{(\ell,j)}$ is a one-dimensional vector associated with the ℓ-th term and coordinate j, and there is no need to work explicitly with $F_{i_1 i_2\ldots i_D}$, which is a D-dimensional tensor with n^D components, For example, if $D = 2$, a SOP basis function with two terms has the form

$$\begin{aligned} & c^1(q_1)g^1(q_2) + c^2(q_1)g^2(q_2) \\ &= \sum_{i_1} f_{i_1}^{(1,1)}\theta_{i_1}^{1}(q_1)\sum_{i_2} f_{i_2}^{(1,2)}\theta_{i_2}^{2}(q_2) \\ &+ \sum_{i_1} f_{i_1}^{(2,1)}\theta_{i_1}^{1}(q_1)\sum_{i_2} f_{i_2}^{(2,2)}\theta_{i_2}^{2}(q_2) \\ &= \sum_{\ell=1}^{2}\sum_{i_1}\sum_{i_2} f_{i_1}^{(\ell,1)} f_{i_2}^{(\ell,2)}\theta_{i_1}^{1}(q_1)\theta_{i_2}^{2}(q_2). \end{aligned} \quad (3)$$

$F_{i_1 i_2\ldots i_D} = \sum_{\ell=1}^{R}\prod_{j=1}^{D} f_{i_j}^{(\ell,j)}$ represents the function in the primitive $\prod_{j=1}^{D}\theta_{i_j}^{j}(q_j)$ basis. This SOP format for multidimensional functions is known as the canonical polyadic (CP) decomposition for tensors in the applied mathematics literature[51–53] (also called parallel factor decomposition or separated representation). Truncating the sum at a given rank R gives a reduced rank approximation for F. The CP format has been successfully applied to the calculation of many-body electronic integrals and wavefunctions.[54–59] Because the factors in the terms are not chosen from a pre-determined set, our basis functions are in the CP format.

There are other reduced (compressed) tensor formats which could be used to compactly represent $F_{i_1 i_2\ldots i_D}$. The most familiar compression of this type, for a 2-D problem, is the singular value decomposition. For $D > 2$, different decompositions exist.[60] In the Tucker format[60–62], $F_{i_1 i_2\ldots i_D} = \sum_{\ell_1=1}^{L_1}\cdots\sum_{\ell_D=1}^{L_D} K_{\ell_1\ell_2\ldots\ell_D}\prod_{j=1}^{D} a_{i_j\ell_j}^{(j)}$. where K is called the core tensor and $L_j < n_j \quad \forall j = 1\ldots D$ and the $a^{(j)}$ are $n_j \times L_j$ matrices. This format is equivalent to the one used by MCTDH.[36] The Hierarchical Tucker format[63–65] (of which the tensor train format[66] is a particular case) is a compromise between the Tucker format and the CP format. It was first introduced by developers of MCTDH[67].

In this article we propose a procedure for making SOP basis functions in the form of Eq. (2). How do we make the basis functions? We shall begin with a function having one term (i.e. with rank 1) that is obtained from $F_{i_1 i_2\ldots i_D} = \prod_{j=1}^{D} f_{i_j}^{(1,j)}$ with some random $f_{i_j}^{(1,j)}$ and obtain basis functions (see the next subsection) by applying the Hamiltonian operator. Throughout this paper we shall assume that the Hamiltonian is also a sum of products,

$$H(q_1,\ldots,q_D) = \sum_{k=1}^{T}\prod_{j=1}^{D} h_{kj}(q_j), \quad (4)$$

where h_{kj} is a one-dimensional operator acting in a Hilbert space associated with coordinate q_j. Kinetic energy operators (KEOs) almost always have this form. If the potential is not in SOP form it can be massaged into SOP form by using, for example, potfit[36,37], multigrid potfit[68], or neural network methods[69–72].

B. Shifted power method

In this subsection we explain how SOP basis functions are made by applying the Hamiltonian. In the $\prod_{j=1}^{D}\theta_{i_j}^{j}(q_j)$ basis the SOP basis functions are represented by the $F_{i_1 i_2 \ldots i_D}$ coefficients in Eq. (2). We use the power method, the simplest iterative method[73,74], to determine the $f_{i_j}^{(\ell,j)}$. Let $\mathbf{F}^{(0)}$ be a random start vector of the form of Eq. (2) and $\mathbf{V}_{E_{\max}}$ be the eigenvector associated with the eigenvalue, $E_{\max}$, whose absolute value is largest. Throughout this paper we shall assume that the minimum potential energy is zero and therefore that all eigenvalues of $\mathbf{H}$, the finite matrix representing the Hamiltonian in the primitive $\prod_{j=1}^{D}\theta_{i_j}^{j}(q_j)$ basis, are positive. In this case, $E_{\max}$ is simply the largest eigenvalue. Assuming $(\mathbf{F}^{(0)})^T\mathbf{V}_{E_{\max}} \neq 0$,

$$\lim_{N_{\rm pow}\to\infty} \mathbf{H}^{N_{\rm pow}}\mathbf{F}^{(0)} \to \mathbf{V}_{E_{\max}} \ . \tag{5}$$

When $N_{\rm pow}$ is large, $\mathbf{F}^{(N_{\rm pow})} = \mathbf{H}^{N_{\rm pow}}\mathbf{F}^{(0)}$ approaches the eigenvector of $\mathbf{H}$ with the largest eigenvalue. If the Hamiltonian is a SOP (Eq. (4)) then $\mathbf{F}^{(N_{\rm pow})}$ has the form of Eq. (2). The convergence of the power method is known to be slow and to depend on gaps between eigenvalues close to $E_{\max}$ and $E_{\max}$.[73,74] The error is approximately proportional to $(E_{sl}/E_{\max})^{N_{\rm pow}}$, where E_{sl} is the second largest eigenvalue.

We could use the $\mathbf{F}^{(N_{\rm pow})}$ sequence to compute the largest eigenvalue of $\mathbf{H}$. Each $\mathbf{F}^{(N_{\rm pow})}$ has the form of Eq. (2) and hence its storage requires little memory. However, we do not want the largest eigenvalue and we wish to compute more than one eigenvalue. From a reasonable estimate of $E_{\max}$ one can obtain the smallest eigenvalue of $\mathbf{H}$, from the linearly shifted operator

$$\widetilde{H} = H - \sigma\mathbb{1} \ , \quad \sigma = E_{\max} \ . \tag{6}$$

When several eigenstates are desired, one uses a block method which begins with a set of B random start vectors. Alternating successive applications of $\widetilde{\mathbf{H}}$ with a modified Gram-Schmidt orthogonalization, we obtains a set of vectors, each of the form of Eq. (2), which converges to the eigenvectors associated with the lowest eigenvalues of $\mathbf{H}$. These are SOP basis vectors. Orthogonalization requires adding vectors which is done by concatenation. This increases the rank. It is not necessary to orthogonalize after every application of $\widetilde{\mathbf{H}}$, instead the orthogonalization can be done only after each set of $N_{\rm ortho}$ matrix-vector products. The convergence of the shifted block method is somewhat less slow than the convergence of the simple power method and now depends on the gaps between the B smallest eigenvalues and the $(B+1)$th smallest eigenvalue of $\mathbf{H}$. Gaps between the B smallest eigenvalues play no role. Degeneracies, within the block, cause no problems. Rather than shifting with $E_{\max}$, it is better to shift with a value slightly larger than the average of $E_{\max}$ and the $(B+1)$th eigenvalue of $\mathbf{H}$.[75] In practice we use,

$$\sigma_{\rm opt} = \frac{E_{\max} + E_{>}}{2} \ , \tag{7}$$

where $E_{>}$ is an upper bound for the $(B+1)$th eigenvalue of $\mathbf{H}$. The desired eigenvalues of $\mathbf{H}$ correspond to the largest eigenvalues of $\widetilde{\mathbf{H}}$. The algorithm can be first applied to $\mathbf{H}$ with $\sigma = 0$ and $B = 1$ to calculate $E_{\max}$. $E_{>}$ is obtained by running a few iterations of the algorithm with $\sigma = E_{\max}$ and a block size $B+1$. The algorithm also works with the non-optimal shift $\sigma = E_{\max}$.

The vectors obtained by successively applying $(\widetilde{\mathbf{H}})$ to a set of B start vectors and orthogonalizing will, if $N_{\rm pow}$ is large enough, approach the matrix of eigenvectors $\mathbf{V} = (\mathbf{V}_1 \ldots \mathbf{V}_B)$. We denote these vectors

$$\boldsymbol{\mathcal{F}} = (\mathbf{F}_1^{(N_{\rm pow})} \ldots \mathbf{F}_B^{(N_{\rm pow})}) \ . \tag{8}$$

$\boldsymbol{\mathcal{F}}$ can also be used as a basis for representing $\mathbf{H}$, to obtain more accurate eigenvalues and eigenvectors. $\boldsymbol{\mathcal{F}}$ is our SOP basis. Even if $N_{\rm pow}$ is not large enough to ensure that $\boldsymbol{\mathcal{F}}$ is a set of eigenvectors, the subspace spanned by the $\boldsymbol{\mathcal{F}}$ set may be sufficient to obtain good approximations for the smallest eigenpairs by projecting into the space, i.e., by computing eigenpairs of the generalized eigenvalue problem, $\mathbf{H}^{(\boldsymbol{\mathcal{F}})}\mathbf{U} = \mathbf{S}\mathbf{U}\mathbf{E}$, where $\mathbf{H}^{(\boldsymbol{\mathcal{F}})} = \boldsymbol{\mathcal{F}}^T\mathbf{H}\boldsymbol{\mathcal{F}}$ and $\mathbf{S} = \boldsymbol{\mathcal{F}}^T\boldsymbol{\mathcal{F}}$. A simple eigenvalue problem would be sufficient to obtain the eigenvectors if the $\boldsymbol{\mathcal{F}}$ basis set were always perfectly orthogonal. However residual non-orthogonality is present, due to a reduction (compression) step that must be introduced into the algorithm (see section II C). This explains why a generalized eigenvalue problem is be used. $\mathbf{S}$ is computed by doing one-dimensional operations,[53] ($\mathbf{F}$ and $\mathbf{F}'$ being two vectors of the $\boldsymbol{\mathcal{F}}$ set),

$$\mathbf{F}^T\mathbf{F}' = \sum_{\ell=1}^{R}\sum_{\ell'=1}^{R'}\prod_{j=1}^{D}(\mathbf{f}^{(\ell,j)})^T\mathbf{f}'^{(\ell',j)} \tag{9}$$

with $(\mathbf{f}^{(\ell,j)})^T\mathbf{f}'^{(\ell',j)} = \sum_{i_j=1}^{n_j} f_{i_j}^{(\ell,j)} f_{i_j}'^{(\ell',j)}$. $\mathbf{H}^{(\boldsymbol{\mathcal{F}})}$ matrix elements are computed similarly, from 1-D operations (Eq. (11) of section II C followed by a scalar product using an approach similar to that of Eq. (9)). It is advantageous to restart, every $N_{\rm diag}$ iterations, with approximate eigenvectors of $\mathbf{H}$ that are columns of $\boldsymbol{\mathcal{F}}\mathbf{U}$. In our programs $N_{\rm diag}$ can be equal to or a multiple of $N_{\rm ortho}$ (see section III and IV). The approximate eigenvectors used to re-start are sums of B different $\mathbf{F}_k^{(N_{\rm pow})}$ vectors

and are obtained by concatenating, which increases the rank to BR if every $\mathbf{F}_k^{(N_{\rm pow})}$ has rank R.

In this section, we specify the SOP basis functions. Thus far it appears that they are obtained from the vectors of Eq. (8). This, however, is not practical. Applying $\mathbf{H}$ to a vector (and even re-starting and orthogonalizing) increases the rank of the vectors (the number of products in the sum) and causes the memory cost to explode. In the next subsection we outline how to obviate this problem by reducing the rank of the SOP basis functions. The combination of an iterative algorithm for making a SOP basis and rank reduction will only work if the basis vectors generated by the iterative algorithm converge to low-rank vectors. If they do, reducing the rank of the vectors will cause little error. If they do not, and even if eigenvectors are low-rank and linear combinations of basis vectors generated by the iterative algorithm, reducing the rank of the vectors will cause significant error. We expect eigenvectors of $\mathbf{H}$ to be low rank and therefore expect it to be possible to reduce the rank of vectors generated by the shifted power method, which approach eigenvectors. The slow convergence of the power method is thus compensated by the advantage of being able to work with low rank vectors.

C. H application and rank reduction

The key step in the block power method is the application of $\mathbf{H}$ to a vector $\mathbf{F}$ to obtain a new vector $\mathbf{F}'$. With T terms in $\mathbf{H}$, the rank of $\mathbf{F}'$ is a factor of T larger than the rank of $\mathbf{F}$. All vectors are represented as

$$F_{i_1 i_2 \ldots i_D} = \sum_{\ell=1}^{R} s_\ell \prod_{j=1}^{D} \tilde{f}_{i_j}^{(\ell,j)} \text{ with } \sum_{i_j}^{n_j} |\tilde{f}_{i_j}^{(\ell,j)}|^2 = 1 \ , \quad (10)$$

where, for each term (ℓ) and each coordinate (j), $\tilde{f}_{i_j}^{(\ell,j)}$ is a normalized 1-D vector, s_ℓ is a normalization coefficient, and n_j is the number of basis functions for coordinate j. Using normalized 1-D vectors allows us to order the different terms in the expansion. This is useful for identifying dominant terms in the sum. $\mathbf{H}$ can be applied to $\mathbf{F}$ by evaluating 1-D matrix-vector products with matrix representations of 1-D operators $\mathbf{h}_{kj}$ in the $\theta_{i_j}^j$ basis, i.e. $(\mathbf{h}_{kj})_{i_j,i'_j} = \langle \theta_{i_j}^j | h_{kj} | \theta_{i'_j}^j \rangle$, and 1-D vectors $\tilde{f}_{i_j}^{(\ell,j)}$,

$$\begin{aligned}
(\mathbf{F}')_{i'_1 i'_2 \cdots i'_D} &= (\mathbf{H}\mathbf{F})_{i'_1 \ldots i'_D} \\
&= \sum_{i_1,i_2,\cdots,i_D} \sum_{k=1}^{T} \prod_{j'=1}^{D} (\mathbf{h}_{kj'})_{i'_{j'} i_{j'}} \sum_{\ell=1}^{R} \prod_{j=1}^{D} s_\ell \tilde{f}_{i_j}^{(\ell,j)} \\
&= \sum_{k=1}^{T} \sum_{\ell=1}^{R} \prod_{j=1}^{D} \sum_{i_j} (\mathbf{h}_{kj})_{i'_j i_j} s_\ell \tilde{f}_{i_j}^{(\ell,j)}. \qquad (11)
\end{aligned}$$

Applying $\mathbf{H}$ to $\mathbf{F}$, with R terms, yields a vector with RT terms. Owing to the fact that everything is done with 1-D matrix-vector products, generating the vector $\mathbf{F}'$ is inexpensive.

If the rank were not reduced after each matrix-vector product, the rank of a vector obtained by applying $\mathbf{H}$ P times to a start vector with R_0 terms would be $T^P R_0$. If T and/or P is large, one would need more, and not less, memory to store the vector than would be required to store n^D components. Table I shows, for $n = 10$, the maximum value of P for which less memory is needed to store a vector obtained by applying $\mathbf{H}$ P times to a start vector with rank one ($R_0 = 1$). This table clearly reveals that rank reduction is imperative.

TABLE I. Maximum number of products $\mathbf{H}$ $\mathbf{F}$ before losing the memory advantage of the CP format if H has T terms, in D dimensions.

$T \setminus D$	3	6	12	20	30
15	2	5	10	17	25
30	2	4	8	13	20
100	1	3	6	10	15
200	1	2	5	8	13
400	1	2	4	7	11

What algorithm is used to reduce the rank and by how much is the rank reduced? To reduce the rank, we replace

$$\begin{aligned}
F^{\rm old}_{i_1 i_2 \ldots i_D} &= \sum_{\ell=1}^{R_{\rm old}} {}^{\rm old}s_\ell \prod_{j=1}^{D} {}^{\rm old}\tilde{f}_{i_j}^{(\ell,j)} \\
\Longrightarrow F^{\rm new}_{i_1 i_2 \ldots i_D} &= \sum_{\ell=1}^{R_{\rm new}} {}^{\rm new}s_\ell \prod_{j=1}^{D} {}^{\rm new}\tilde{f}_{i_j}^{(\ell,j)} \ , \quad (12)
\end{aligned}$$

where $R_{\rm new} < R_{\rm old}$ and choose ${}^{\rm new}\tilde{f}_{i_j}^{(\ell,j)}$ to minimize $\| \mathbf{F}^{\rm new} - \mathbf{F}^{\rm old} \|$. Making this replacement changes a vector generated by the power method, but because energy levels are computed by projecting into the space spanned by $\boldsymbol{\mathcal{F}}$, numerically exact results can still be obtained. If $R_{\rm new} \sim R_{\rm old}$, $\| \mathbf{F}^{\rm new} - \mathbf{F}^{\rm old} \|$ is small but the memory cost is large. One might choose $R_{\rm new}$, for each reduction, so that $\| \mathbf{F}^{\rm new} - \mathbf{F}^{\rm old} \|$ is less than some threshold. Instead, we use the same $R_{\rm new}$ for all reductions and choose a value small enough that the memory cost is much less than the cost of storing n^D components but large enough that good results are obtained from a relatively small value of $N_{\rm pow}$. Rank reduction is motivated by the realization that when the Hamiltonian is separable, i.e. $H(q_1, \ldots, q_D) = h_1(q_1) + h_2(q_2) + \cdots + h_D(q_D)$, the wavefunctions are all of rank one and when coupling is not huge the rank of wavefunctions is small (it is important to understand that the rank of a wavefunction is not the same as the number of $\prod_{j=1}^{D} \theta_{i_j}^j(q_j)$ basis functions which contribute to it). In general, the stronger the coupling, the larger the required value of $R_{\rm new}$. Note that wavefunctions are represented as linear combinations of basis functions with rank $R_{\rm new}$ and may therefore have rank larger than $R_{\rm new}$.

We use an alternating least squares (ALS) algorithm described in Ref. 53 to determine the ${}^{\rm new}\tilde{f}_{i_j}^{(\ell,j)}$ by minimizing $\parallel \mathbf{F}^{\rm new} - \mathbf{F}^{\rm old} \parallel$. The reduction algorithm needs start values for ${}^{\rm new}\tilde{f}_{i_j}^{(\ell,j)}$. We use the $R_{\rm new}$ terms in $\mathbf{F}^{\rm old}$ with the largest s_ℓ coefficients. Another possibility is to use random start vectors. For all l values ${}^{\rm new}\tilde{f}_{i_j}^{(\ell,j)}$ factors are varied, for a single $j = k$, keeping all the other factors $\tilde{f}_{i_j}^{(\ell,j)} \quad \forall j \neq k$ fixed (and then the vectors are normalized by changing s_ℓ). For each coordinate we solve

$$\frac{\partial \parallel \mathbf{F}^{\rm new} - \mathbf{F}^{\rm old} \parallel}{\partial {}^{\rm new}\tilde{f}_{i_k}^{(\ell,k)}} = 0 \quad \forall \ell, \quad \forall i_k. \tag{13}$$

For a single coordinate k, this requires solving linear systems with an $(R_{\rm new} \times R_{\rm new})$ matrix whose elements are

$$B(\hat{\ell},\tilde{\ell}) = \prod_{\substack{j=1 \\ j\neq k}}^{D} \left({}^{\rm new}\tilde{\mathbf{f}}^{(\hat{\ell},j)} \right)^T {}^{\rm new}\tilde{\mathbf{f}}^{(\tilde{\ell},j)} \tag{14}$$

and with n_k different right-hand-sides, the i_kth of which is

$$d_{i_k}(\hat{\ell}) = \sum_{\ell=1}^{R_{\rm old}} {}^{\rm old}s_l \; {}^{\rm old}\tilde{f}_{i_k}^{(\ell,k)} \prod_{\substack{j=1 \\ j\neq k}}^{D} \left({}^{\rm old}\tilde{\mathbf{f}}^{(\ell,j)} \right)^T \; {}^{\rm new}\tilde{\mathbf{f}}^{(\hat{\ell},j)} \; . \tag{15}$$

${}^{\rm new}\tilde{f}_{i_j}^{(\tilde{\ell},j)}$ is obtained by solving,

$$\sum_{\tilde{\ell}} B(\hat{\ell},\tilde{\ell}) {}^{\rm new}\tilde{f}_{i_j}^{(\tilde{\ell},j)} \;=\; d_{i_j}(\hat{\ell}) \; . \tag{16}$$

Ill-conditioning is avoided with a penalty term as described in Ref. 53. Repeating this for all D coordinates constitutes one ALS iteration, with a computational cost of

$$\mathcal{O}\left(D(R_{\rm new}^3 + n(R_{\rm new}^2 + R_{\rm new}R_{\rm old})) \right), \tag{17}$$

where n is a representative value of n_j. $\mathcal{O}(DnR_{\rm new}^2)$ is the cost of making the $\mathbf{B}$ matrices ($\mathbf{B}$ matrices for successive coordinates are made by updating[53]), $\mathcal{O}(DnR_{\rm new}R_{\rm old})$ is the cost of computing the right-hand-sides, and $\mathcal{O}(DR_{\rm new}^3)$ is the cost of solving the linear systems.

One could iterate the ALS algorithm until $\parallel \mathbf{F}^{\rm new} - \mathbf{F}^{\rm old} \parallel$ is less than some pre-determined threshold. Instead, we fix the number of ALS iterations, N_{ALS}, on the basis of preliminary tests. We do this because computing $\parallel \mathbf{F}^{\rm new} - \mathbf{F}^{\rm old} \parallel$ is costly as it requires the calculation of many scalar products with vectors of rank $(R_{\rm new} + R_{\rm old})$, which scales as $\mathcal{O}((R_{\rm new} + R_{\rm old})^2)$. If $\mathbf{F}^{\rm new}$ determined by fixing the number of ALS iterations is not a good approximation to $\mathbf{F}^{\rm old}$ then we alter the vector obtained from the block power method more than we would like, however, this changes only a basis vector and does not preclude computing accurate energy levels. More effective or more efficient reduction algorithms exist, such as the Newton method of Ref. 51 and 76 and the conjugate gradient method of Ref. 77, but we have not tried to use them.

D. Combination of block power method with rank reduction: Reduced Rank Block Power Method (RRBPM)

Three operations in the block power method cause the rank of the basis vectors to increase and must therefore be followed by rank reduction. As already discussed, applying $\mathbf{H}$ to a vector increases its rank by a factor of T. Orthogonalization requires adding vectors and the rank of the sum of two vectors is the sum of their ranks. After solving the generalized eigenvalue problem, the basis vectors are updated by replacing them with linear combinations (the coefficients being elements of the eigenvector matrix) of basis vectors; this also increases the rank. Orthogonalization and updating increase the rank by much less than applying $\mathbf{H}$, nevertheless if they are not followed by a rank reduction the rank of the basis vectors will steadily increase during the calculation.

The algorithm we use is:

1. Define B random rank-one initial vectors $\mathbf{F}_b$ with elements $\prod_{j=1}^{D} f_{b,i_j}^{(1,j)}$ for $b = 1,\ldots,B$, $i_j = 1,\ldots,n_j$.
2. Orthogonalize the $\mathbf{F}_b$ set with a modified Gram-Schmidt procedure adapted to the SOP structure.
3. First reduction step: if $B > r$, reduce the rank of the orthogonalized $\mathbf{F}_b$ to r using ALS.
4. Iterate:
   (a) Apply $\mathbf{F}_b \leftarrow (\mathbf{H} - \sigma \mathbb{1})\mathbf{F}_b \;\; \forall b = 1,\ldots,B$.
   (b) Main reduction step: reduce all the $\mathbf{F}_b$ to rank r using ALS.
   (c) Every $N_{\rm ortho}$ iterations:
      i. Orthogonalize the $\{\mathbf{F}_b\}$ set using a SOP-adapted modified Gram-Schmidt procedure.
      ii. Reduce the rank of all the $\mathbf{F}_b$ to r using ALS.
   (d) Every $N_{\rm diag}$ iterations (multiple of $N_{\rm ortho}$):
      i. Compute $\mathbf{H}_{bb'}^{(\boldsymbol{\mathcal{F}})} = \mathbf{F}_b^T \mathbf{H} \mathbf{F}_{b'}$ and $\mathbf{S}_{bb'} = \mathbf{F}_b^T \mathbf{F}_{b'}$.
      ii. Solve the generalized eigenvalue problem $\mathbf{H}^{(\boldsymbol{\mathcal{F}})}\mathbf{U} = \mathbf{SUE}$ where $\mathbf{E}$ is the diagonal matrix of eigenvalues and $\mathbf{U}$ is the matrix of eigenvectors.
      iii. Update the vectors $\mathbf{F}_{b'}^{\rm new} = \sum_{b=1}^{B} \mathbf{U}_{bb'}\mathbf{F}_b$.
      iv. $\mathbf{F}_b \leftarrow$ reduction of $\mathbf{F}_b^{\rm new}$ to rank r using ALS, $\forall b = 1,\ldots,B$.

At step 2 and step 4(c)i, we orthogonalize with a Gram-Schmidt procedure adapted to exploit the SOP structure of the vectors. This requires computing scalar products of 1-D $\tilde{\mathbf{f}}_b^{(1,j)}$ vectors and adding $\mathbf{F}_b$ vectors. At step 4(d)ii we also add vectors. Adding vectors is done by concatenation.

The program that implements the algorithm is not fully optimized, but we have parallelized some of its steps. Step 4,a is embarrassingly parallel because $(H-\sigma\mathbb{1})$ can be applied to each vector separately. In all the orthogonalization steps, there are two nested loops on indices b and b':

```
for b = 1 to B
    F_b ← F_b/‖F_b‖
    for b' = b + 1 to B
        F_b' ← F_b' − (F_b^T F_b')F_b
    end
end
```

The internal loop over b' is parallelized. The reduction of steps 3, 4(b), 4(c) ii, and 4(d)iv is also done in parallel.

After step 2, the maximum rank of a basis vector is B. The first reduction step 3 is only done if $B > r$. After the application of $(H-\sigma\mathbb{1})$, every basis vector has rank Tr (or TB during the first passage if $B < r$), which is reduced to r in step 4(b). After the orthogonalization step 4(c)i, as well as after the diagonalization step 4(d)ii the rank is Br, but is immediately reduced to r.

The memory cost is the memory required to store the largest rank vectors. If $T > B$, then the largest rank vectors are those obtained after application of $(H-\sigma\mathbb{1})$ and they have rank of Tr. If $B > T$, then the largest rank vectors are those obtained after the orthogonalization and they have rank of Br. Therefore, the memory cost scales as

$$\begin{aligned} &\mathcal{O}(nDBTr) \text{ if } T > B \\ &\mathcal{O}(nDB^2r) \text{ if } T < B\ . \end{aligned} \tag{18}$$

The CPU cost is dominated by the reduction steps (scalings are given in section II C). The cost of one matrix vector product scales as $\mathcal{O}(TRDn^2)$ (see Eq. (11)).

III. BILINEARLY COUPLED HARMONIC OSCILLATORS

We first test the rank-reduced block power method (RRBPM) by computing eigenvalues of a simple Hamiltonian for which exact energy levels are known,

$$H(q_1,\ldots,q_D) = \sum_{j=1}^{D} \frac{\omega_j}{2}\left(p_j^2+q_j^2\right) + \sum_{\substack{i,j=1\\ i>j}}^{D} \alpha_{ij} q_i q_j \tag{19}$$

with $p_j = -\imath\frac{\partial}{\partial q_j}$. The product basis $\prod_{j=1}^{D}\theta_{i_j}^{j}(q_j)$ is made from n_j harmonic oscillator basis functions for each coordinate, i.e. eigenfunctions of $\frac{\omega_j}{2}\left(p_j^2+q_j^2\right)$. The exact levels of H, obtained by transforming to normal coordinates are

$$E_{m_1,m_2,\cdots,m_D} = \sum_{j=1}^{D}\nu_j\left(\frac{1}{2}+m_j\right), \text{ with } m_j = 0,1,\cdots. \tag{20}$$

The normal mode frequencies ν_j are square roots of the eigenvalues of the matrix $\mathbf{A}$ whose elements are $A_{ii}=\omega_i^2$ and $A_{ij} = \alpha_{ij}\sqrt{\omega_i}\sqrt{\omega_j}$. Despite the simplicity of the Hamiltonian, it is a good test of the RRBPM. Our goal is to determine whether it is possible to obtain accurate levels when the coupling is large enough to significantly shift levels. Is it possible to compute accurate levels with a value of r small enough the memory cost of the RRBPM is significantly less than the memory cost of a method that requires storing all n^D components of vectors? Does the ALS procedure for rank reduction work well in these conditions?

A. 6-D coupled oscillators

We arbitrarily choose the coefficients of Eq. (19),

$$\omega_j = \sqrt{j/2},\ j = 1,\ldots,6. \tag{21}$$

For simplicity, the same value $\alpha_{ij} = 0.1$ is given to all the coupling constants. The coefficient of a quadratic coupling term is about 14% of the smallest ω_j. The coupling shifts the frequencies of transitions $0 \to 9$ and $0 \to 39$ by about 2 percent. For realistic transitions around 3000 cm^{-1}, this would correspond to a shift of 60 cm^{-1}; the coupling is therefore significant. Energies computed with the parameters of table II are reported in table III. For a given choice of the Hamiltonian parameters and the basis size parameters (n_j), one expects the accuracy to be limited by the values of r, N_{ALS}, B, and the maximum value of N_{pow}. Increasing any of these will increase the accuracy. Regardless of the values of r and N_{ALS}, accurate energies can be obtained by increasing B. The r, N_{ALS}, B, and $\text{Max}(N_{\text{pow}})$ values in table II were determined by testing various values, but many sets of parameter values work well. For the bilinearly coupled Hamiltonian, the largest eigenvalue of $\mathbf{H}$ could be estimated from the largest diagonal matrix element, but we compute it with the power method (no shift, block size of one). We choose a shift close to the largest eigenvalue. This shift is not the optimal value given in Eq. (7). Decreasing the shift, as explained in section II B, slightly accelerates the convergence. For 3000 iterations the calculation requires approximately 6 min on a computer with 2 Quad-Core AMD 2.7 GHz processors, using all of the processing cores.

Two versions of the algorithm are tested, with and without updating eigenvectors (item 4d in the algorithm of section II D). When there is no updating during the iterations, the subspace diagonalization is done from time to time, only to follow the convergence of the results.

TABLE II. Parameters for the calculation with $D = 6$.

D	6
ω_j	$\sqrt{j/2}$
α_{ij}	0.1
$n_j, \quad \forall j = 1, \ldots, 6$	10
Reduction rank r	10
N_{ALS}	10
Block size B	40
Maximum N_{pow}	3000
E_{max} estimate	80.36
Energy shift σ	81
N_{ortho}	20
N_{diag}	20

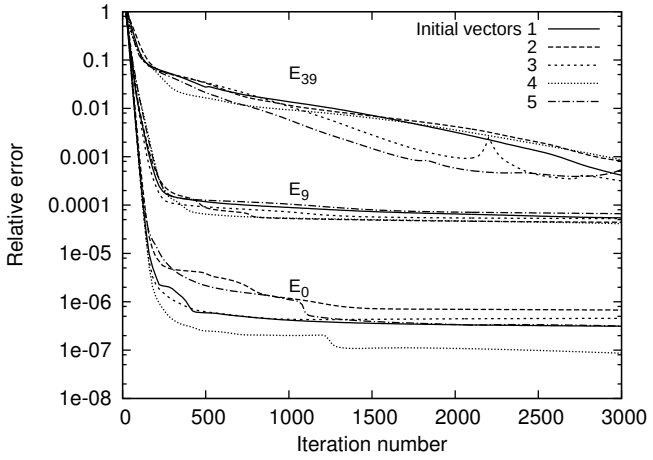

FIG. 1. Convergence curves for the lowest eigenvalue E_0, the highest eigenvalue of the block E_{39} and another one inside the block, E_9. The y axis is the relative error (logarithmic scale) and the x axis is the iteration number, N_{pow}. Results are shown for five calculations using five different randomly chosen sets of start vectors.

With the parameters of table III and implementing updating, the zero-point energy (ZPE) E_0 is recovered with a 10^{-7} relative error (6 significant digits). For low-lying levels there is no significant difference between energies obtained with and without updating. As expected, the lowest eigenvalues are the most accurately determined, with a relative error of less than 10^{-4} for the first ten eigenvalues. For the higher eigenvalues in the block, a little less than one order of magnitude is gained by updating. The quality of the highest eigenvalues of the block is less good, but even the largest calculated eigenvalue E_{39} has a relative error of order 10^{-4}, despite the very low rank that we have chosen ($r = 10$) and its proximity with neighboring eigenvalues. These results confirm that even with basis functions with only $r = 10$ terms, good accuracy can be achieved. With these parameters, the calculation requires only 20 MB of memory.

To verify that accuracy is not limited by N_{pow} we plotted, in Fig. 1, the relative error for three eigenvalues as a function of N_{pow}, with five different start vectors. We

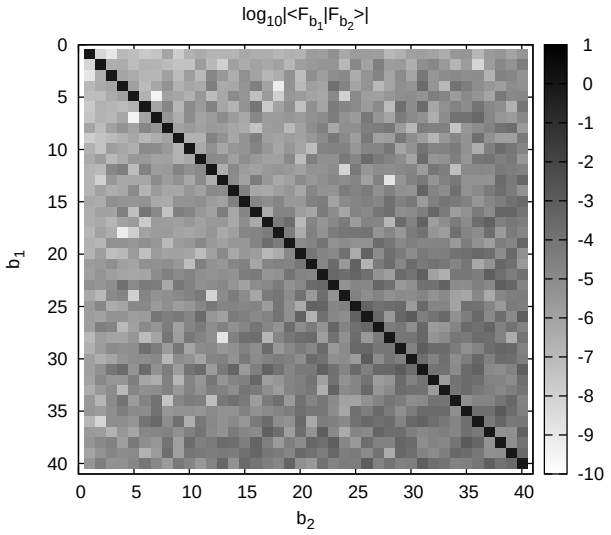

FIG. 2. Grey scale logarithmic representation of the overlap matrix, $\log_{10} |\mathbf{S}_{bb'}| = \log_{10} |\mathbf{F}_b^T \mathbf{F}_{b'}|$ after the last (orthogonalization + reduction) step. The black pixels correspond to approximately 1.

focus attention on the ZPE, E_0, the highest state of the block, E_{39} and an intermediate eigenvalue, E_9. As expected, the lowest eigenvalue converges faster than E_9 and E_{39}. Error decreases rapidly for the first 250 iterations and thereafter more slowly. The results depend little on the choice of the start vector (maximum variation of about an order of magnitude after 3000 iterations).

As explained in the last section, it is crucially important to reduce the rank of the SOP vectors. However, the reduction changes the vectors and the space they span. To assess how much the functions are changed by reduction, we examine the overlap matrix $\mathbf{S}$ after the reduction following orthogonalization. The overlap matrix after iteration number 3000 is shown in Fig. 2. The reduction reduces the rank from Br to r. Reduction causes this matrix to differ from an identity matrix. Each pixel corresponds to one element of the overlap matrix $S_{bb'}$ with a logarithmic grey-scale. After normalizing the diagonal elements to one, the largest off-diagonal elements are of order 10^{-3}. Neglecting the off-diagonal elements creates sufficient errors in the spectrum to justify the choice of a generalized eigenvalue algorithm in the algorithm described in section II D.

Energy errors in this section are differences between RRBPM energies and exact (from normal frequencies) energies. In order to test the ability of the RRBPM to determine eigenvalues of a basis representation of the Hamiltonian operator, it would be better to take differences of RRBPM energies and exact eigenvalues of $\mathbf{H}$. By computing exact eigenvalues of $\mathbf{H}$ (using a product basis Lanczos method), we have verified that the accuracy of RRMP energies in not limited by the primitive basis. The Lanczos eigenvalues agree with the exact en-

TABLE III. Energy levels of the 6D coupled oscillator Hamiltonian. From left to right: energy label n, exact energy, RRBPM energy with no updating, relative error, RRBPM energy with updating, relative error. The last column is the normal mode assignment. Parameters from table II are used.

n	$E_{n,\mathrm{th}}$	$E^{(1)}_{n,\mathrm{num}}$ (no update)	$\frac{E^{(1)}_{n,\mathrm{num}}-E_{n,\mathrm{th}}}{E_{n,\mathrm{th}}}$	$E^{(2)}_{n,\mathrm{num}}$ (with update)	$\frac{E^{(2)}_{n,\mathrm{num}}-E_{n,\mathrm{th}}}{E_{n,\mathrm{th}}}$	Assignment
0	3.8164041	3.8164063	6.0×10^{-7}	3.8164053	3.1×10^{-7}	-
1	4.5039223	4.5039260	8.2×10^{-7}	4.5039269	1.0×10^{-6}	ν_1
2	4.7989006	4.7989650	1.3×10^{-5}	4.7989351	7.2×10^{-6}	ν_2
3	5.0262787	5.0262979	3.8×10^{-6}	5.0263184	7.9×10^{-6}	ν_3
4	5.1914405	5.1914748	6.6×10^{-6}	5.1914768	7.0×10^{-6}	$2\,\nu_1$
5	5.2196500	5.2196889	7.4×10^{-6}	5.2197161	1.3×10^{-5}	ν_4
6	5.3938508	5.3940540	3.8×10^{-5}	5.3939608	2.0×10^{-5}	ν_5
7	5.4864188	5.4865101	1.7×10^{-5}	5.4864995	1.5×10^{-5}	$\nu_1+\nu_2$
8	5.5886301	5.5886564	4.7×10^{-6}	5.5887086	1.4×10^{-5}	ν_6
9	5.7137969	5.7139860	3.3×10^{-5}	5.7141074	5.4×10^{-5}	$\nu_2+\nu_3$
⋮	⋮	⋮	⋮	⋮	⋮	⋮
19	6.3763473	6.3797375	5.3×10^{-4}	6.3770751	1.1×10^{-4}	$\nu_2+\nu_5$
20	6.4013151	6.4041135	4.4×10^{-4}	6.4017330	6.5×10^{-5}	$2\,\nu_1+\nu_3$
21	6.4295246	6.4331587	5.7×10^{-4}	6.4305409	1.6×10^{-4}	$\nu_3+\nu_4$
22	6.4689153	6.4728294	6.1×10^{-4}	6.4693700	7.0×10^{-5}	$\nu_1+2\,\nu_2$
23	6.5664770	6.5686726	3.3×10^{-4}	6.5665373	9.2×10^{-6}	$4\,\nu_1$
24	6.5711267	6.5731186	3.0×10^{-4}	6.5720287	1.4×10^{-4}	$\nu_2+\nu_6$
⋮	⋮	⋮	⋮	⋮	⋮	⋮
34	6.8896648	6.9025957	1.9×10^{-3}	6.8924037	4.0×10^{-4}	$\nu_1+\nu_2+\nu_4$
35	6.9236715	6.9354401	1.7×10^{-3}	6.9253179	2.4×10^{-4}	$\nu_1+2\,\nu_3$
36	6.9636666	6.9665546	4.1×10^{-4}	6.9645821	1.3×10^{-4}	$2\,\nu_1+\nu_6$
37	6.9712976	6.9796692	1.2×10^{-3}	6.9721608	1.2×10^{-4}	$2\,\nu_5$
38	6.9912717	6.9990297	1.1×10^{-3}	6.9928003	2.2×10^{-4}	$2\,\nu_2+\nu_3$
39	6.9918761	7.0057516	2.0×10^{-3}	6.9947683	4.1×10^{-4}	$\nu_4+\nu_6$

ergies to within 10^{-10}. The Lanczos calculation is easy in 6D.

To understand how changing the parameters in table II affects energy levels, we did a series of calculations keeping all but one of the parameters fixed at the values in table II. These tests are done for a single start vector. In Fig. 3(a) we show how the relative error, for three energies, decreases as r is increased. Energies are, of course, more accurate when the reduction rank is larger. However, the error vs r curves flatten as r increases. Even with relatively small values of r, good accuracy is obtained, e.g., the relative error of the ZPE is 3×10^{-5} with $r=5$. The most costly part of the calculation is the reduction which involves a solution of linear equations scaling as r^3. Fig. 3(b) shows that although the error increases as $N_{\mathrm{diag}}=N_{\mathrm{ortho}}$ is increased, reorthogonalizing and updating less frequently does not significantly degrade the accuracy. Starting at about $N_{\mathrm{diag}}>100$, the error does increase somewhat. The ZPE is accurate even with a large N_{diag} (without any orthogonalization, all the basis vectors would tend to the lowest eigenstate). It is fortunate that $N_{\mathrm{diag}}=N_{\mathrm{ortho}}$ need not be large because except for updating, everything can be parallelized over the vectors in the block. The subspace diagonalization requires $\mathcal{O}(B^2)$ scalar products to compute the matrices and $\mathcal{O}(B^3)$ operations for the direct diagonalization. However these operations are not costly because they are repeatedly infrequently. Another important parameter is the number of ALS iterations (see section II C). We have used ALS to reduce the rank of test functions and observed that the error vs number of ALS iterations curve often flattens and then decreases slowly (not shown). The convergence behavior also depends on the vector whose rank is being reduced. It is therefore extremely likely that some of the reduced rank vectors we use as basis vectors in the RRBPM, where we fix the number of ALS iterations, have significant errors. The fact that the basis vectors differ from those one would have with a standard block power method is not necessarily a cause for concern because we project into the space spanned by the basis vectors. This works for the same reason that it is possible to use inexact spectral transforms.[14,78–81] We have nevertheless done tests with several values of N_{ALS}. With $N_{ALS}=10$ accurate energies are obtained. If N_{ALS} is too small, errors are larger and the errors of higher states are larger. None of the 39 levels we compute are significantly improved by doubling N_{ALS}. Replacing $\sigma=81$ with $\sigma=44$, consistent with the optimal value of Eq. (7) slightly accelerates the convergence. Increasing the block size B accelerates convergence, but the effect is most important when the iteration number is less than about 500.

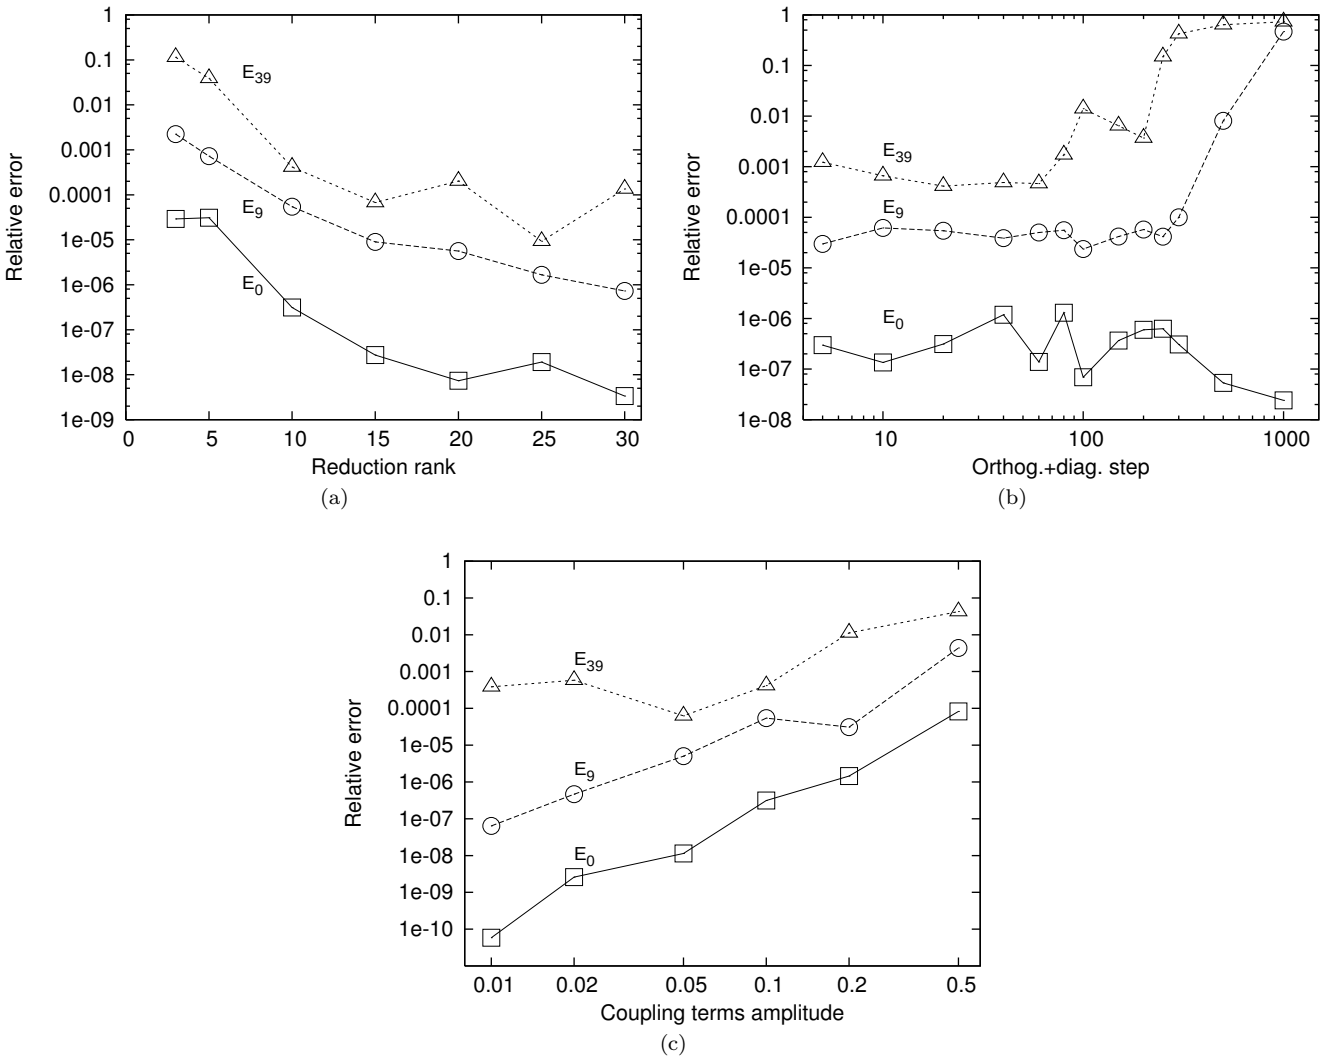

FIG. 3. Relative error of E_0, E_9 and E_{39} after Max($N_{\rm pow}$)=3000 RRBPM iterations, on a logarithmic scale and as a function of: (a) the reduction rank r; (b) the orthogonalization-diagonalization step size (here $N_{\rm ortho} = N_{\rm diag}$); (c) the coupling term amplitude α_{ij}.

Of course as the coupling α_{ij} is increased the accuracy of a method based on rank reduction must degrade. The α_{ij} value we have used forces a realistic mixing. Although low-rank approximations for the lowest states are good, many of the $\prod_{j=1}^{D} \theta_{i_j}^{j}(q_j)$ basis functions are mixed. Nevertheless, to force the RRBPM to fail, as it must, we increased α_{ij}. The results are displayed in Fig. 3(c). All other parameters have the values in table II except the energy shift which has been adapted for each run because the spectral range changes. As expected, increasing the coupling increases the error. When α_{ij} is larger than 0.1 better accuracy can be obtained by increasing r.

B. 20-D coupled oscillators

The basis used in the 6D calculation of the previous subsection has 10^6 functions. Obviously, direct diagonalisation methods cannot be used, but it is easy compute eigenvalues of $\mathbf{H}$ with the Lanczos algorithm. In this subsection, we present results for a 20D calculation with the Hamiltonian of Eq. (19) and the parameters listed in table IV. In this case, the Lanczos calculation is also impossible. With $n_j = 10$, a single vector has 10^{20} components and to keep a vector in memory one would need $\sim 8\times10^{11}$ GB. Using the RRBPM and less than 1 GB, we compute about 50 energy levels with relative accuracy of about 10^{-5}. The Hamiltonian has 210 terms. The shift value is slightly larger than half the maximum eigenvalue ($\simeq 568$). The shift value is an optimal shift (Eq. (7)), obtained from $E_{\rm max}$ and $E_>$. To determine $E_>$, a few it-

erations of the RRBPM are done using the non-optimal shift $\sigma \simeq E_{\max}$ and a block of size $B+1$.

TABLE IV. Numerical parameters for the calculations with $D = 20$ coupled oscillators.

D	20
ω_j	$\sqrt{j/2}$
α_{ij}	0.1
$n_j, \quad \forall j = 1, \ldots, 20$	10
Reduction rank r	20
N_{ALS}	10
Block size B	56
Maximum N_{pow}	5000
$E_{\max}$ estimate	568
Energy shift σ	320
N_{ortho}	40
N_{diag}	40

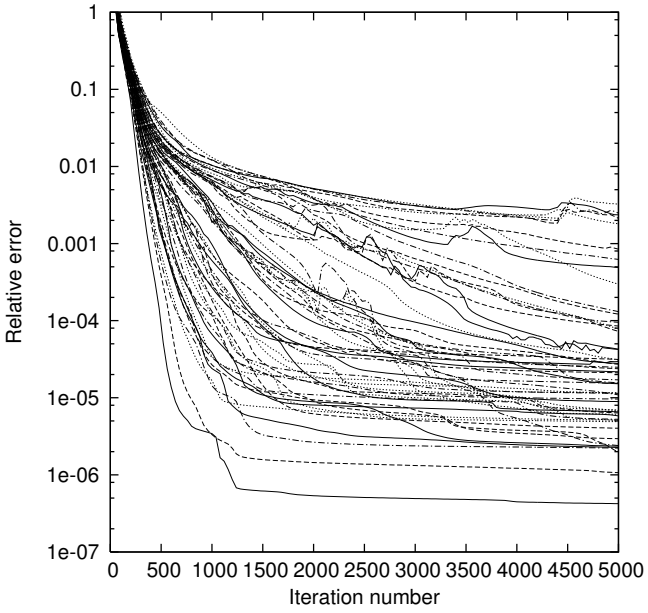

FIG. 4. Convergence curves of the energy levels of the 20-D coupled oscillator Hamiltonian. Parameters of table IV are used. The y-axis is the relative error on a logarithmic scale and the x-axis is the iteration number N_{pow}.

Some computed eigenvalues are presented in table V along with the corresponding exact eigenvalues calculated from the normal mode frequencies. Energies are assigned with normal mode frequencies. In this case it is not possible to compare the RRBPM results with those obtained from a standard Lanczos calculation due to the memory cost of the latter. We therefore have no choice but to calculate relative errors using exact eigenvalues, assuming that the finite basis introduces no significant error, as shown in the 6D case. The agreement between theoretical and numerical eigenvalues is good, with errors of the order of 10^{-5}, and especially good for the lowest energies.

TABLE V. Energy levels of the 20 coupled oscillator Hamiltonian of Eq. (19). From left to right: energy level number, exact energy level, RRBPM energy level, relative error, normal mode assignment.

n	$E_{n,\mathrm{th}}$	$E_{n,\mathrm{num}}$	$\frac{E_{n,\mathrm{num}}-E_{n,\mathrm{th}}}{E_{n,\mathrm{th}}}$	Assignment
0	21.719578	21.719587	4.2×10^{-7}	-
1	22.398270	22.398294	1.1×10^{-6}	ν_1
2	22.691775	22.691826	2.2×10^{-6}	ν_2
3	22.917012	22.917129	5.1×10^{-6}	ν_3
4	23.076962	23.077014	2.3×10^{-6}	$2\,\nu_1$
5	23.106960	23.107006	2.0×10^{-6}	ν_4
6	23.274380	23.274502	5.3×10^{-6}	ν_5
7	23.370467	23.370629	6.9×10^{-6}	$\nu_1+\nu_2$
8	23.425814	23.425951	5.8×10^{-6}	ν_6
9	23.565153	23.565222	2.9×10^{-6}	ν_7
⋮	⋮	⋮	⋮	⋮
20	24.049160	24.049374	8.9×10^{-6}	$2\,\nu_1+\nu_2$
21	24.079158	24.079914	3.1×10^{-5}	$\nu_3+\nu_4$
22	24.104506	24.104878	1.6×10^{-5}	$\nu_1+\nu_6$
23	24.114446	24.114570	5.2×10^{-6}	$2\,\nu_3$
⋮	⋮	⋮	⋮	⋮
30	24.342665	24.343080	1.7×10^{-5}	$\nu_1+2\,\nu_2$
31	24.346217	24.346365	6.1×10^{-6}	ν_{14}
32	24.373625	24.373996	1.5×10^{-5}	$\nu_1+\nu_8$
33	24.398012	24.398676	2.7×10^{-5}	$\nu_2+\nu_6$
⋮	⋮	⋮	⋮	⋮
40	24.532333	24.533376	4.3×10^{-5}	ν_{16}
41	24.537351	24.539130	7.3×10^{-5}	$\nu_2+\nu_7$
42	24.567902	24.570246	9.5×10^{-5}	$\nu_1+\nu_2+\nu_3$
43	24.611100	24.613013	7.8×10^{-5}	$\nu_1+\nu_{10}$
⋮	⋮	⋮	⋮	⋮
50	24.709939	24.725314	6.2×10^{-4}	ν_{18}
51	24.721068	24.765667	1.8×10^{-3}	$\nu_1+\nu_{11}$
52	24.727852	24.786084	2.4×10^{-3}	$3\,\nu_1+\nu_2$
53	24.757850	24.810515	2.1×10^{-3}	$\nu_1+\nu_2+\nu_4$
54	24.762587	24.823982	2.5×10^{-3}	$\nu_3+\nu_7$
55	24.783198	24.863299	3.2×10^{-3}	$2\,\nu_1+\nu_6$

Convergence curves are presented in Fig. 4, which shows the relative error for all 56 eigenvalues of a $B \times B$ block as a function of the iteration number. The error does not decrease monotonically because rank reduction can worsen the basis. The smallest energies are most accurate. The largest energies are not converged with respect to the number of iterations even with a maximum value of N_{pow} equal to 5000. Most of the eigenvalues stabilize with an error less than about 10^{-4}. Although this figure is similar to Fig. 1, convergence is slower. Increasing the number of coordinates from 6 to 20 increases the spectral range of $\mathbf{H}$ and slows the convergence of the block power method. It also makes wavefunctions more complex and reduces the accuracy of low-rank vectors. Ideally, the cost of the calculation would scale linearly with D. In practice, linear scaling is not realized because increasing the number of coupled degrees of free-

dom increases the rank required to represent wavefunctions. This calculation, with 5000 iterations takes 2.5 days using 14 processors. The memory cost is a factor of 50 (consistent with Eq. (18)) greater than in the 6D case because the 20D Hamiltonian has many more terms than the 6D Hamiltonian, but is still less than 1 GB.

IV. APPLICATION TO A 12-D MODEL FOR ACETONITRILE (CH_3CN)

It is extremely encouraging that with the RRBPM it is possible to calculate accurate energies of a 20D Hamiltonian. The Hamiltonian of the previous section is chosen to facilitate testing the numerical approach. Does the RRBPM also work well for a realistic Hamiltonian? One might worry that rank reduction will make it impossible to compute accurate energies for a Hamiltonian with a large number of coupling terms with realistic magnitudes. In this section we confirm that it works well when applied to a 12D Hamiltonian in normal coordinates with a quartic potential. The potential is for acetonitrile CH_3CN.

The normal coordinates are labelled q_k with $k = 1, 2, \cdots 12$. $q_k,\ k = 5, 6, \cdots$ are two-fold degenerate (they correspond to $q_5 \ldots q_8$ of Ref. 82). We use the $J = 0$ normal coordinate kinetic energy operator (KEO), but omit the $\pi - \pi$ cross terms and the potential-like term,[83]

$$K = -\frac{1}{2}\sum_k \omega_k \frac{\partial^2}{\partial q_k^2}. \tag{22}$$

We use the quartic potential of Ref. 50 which is inferred from Ref. 82. Ref. 82 gives force constants calculated with a hybrid coupled cluster/density functional theory method. They also report vibrational levels computed from subspaces determined with second order perturbation theory. The potential is

$$V(q_1, \ldots, q_{12}) = \frac{1}{2}\sum_{i=1}^{12}\omega_i q_i^2 + \frac{1}{6}\sum_{i=1}^{12}\sum_{j=1}^{12}\sum_{k=1}^{12}\phi^{(3)}_{ijk} q_i q_j q_k + \frac{1}{24}\sum_{i=1}^{12}\sum_{j=1}^{12}\sum_{k=1}^{12}\sum_{\ell=1}^{12}\phi^{(4)}_{ijk\ell} q_i q_j q_k q_\ell \ . \tag{23}$$

According to Ref. 82, constants smaller than 6 cm^{-1} were not reported. The force constants must satisfy symmetry relations given by Henry and Amat[84,85], but usually expressed in terms of the Nielsen k constants[86] that correspond to the ϕ in Eq. (23). All of the (937) non-zero ϕ constants can be determined from (358) constants that Henry and Amat denote k^0, k^1, and k^2. Ref. 82 reports 132 ϕ. Some of the missing ϕ are less than 6 cm^{-1}, some of the missing ϕ are not small and can be determined from those reported, some of the missing ϕ are not small and cannot be determined from those reported. Avila et al. assume that the force constants reported in Ref. 82 are the ϕ that correspond to the k^0 force constants (the corresponding ϕ are derivatives of the potential with respect to the x components of the doubly degenerate normal coordinates) and, because they do not have values for them, put the k^1 and k^2 force constants equal to zero. The resulting potential is invariant with respect to the C_{3v} operations. There are 299 (108 cubic and 191 quartic) coupling terms in the potential. Poirier and Halverson infer a different potential from the force constants published in Ref. 82.[87] Most energy levels on their potential differ from their counterparts on the potential of Ref. 50 by less than 1 cm^{-1}. Either potential could be used to test the RRBPM, but we choose the potential of Ref. 50.

To make the $\prod_{j=1}^{D}\theta^j_{i_j}(q_j)$ basis we use the harmonic oscillator basis functions that are eigenfunctions of the quadratic part of the Hamiltonian. The harmonic frequencies are[82] (in cm^{-1}) $\omega_1 = 3065$, $\omega_2 = 2297$, $\omega_3 = 1413$, $\omega_4 = 920$, $\omega_5 = \omega_6 = 3149$, $\omega_7 = \omega_8 = 1487$, $\omega_9 = \omega_{10} = 1061$, $\omega_{11} = \omega_{12} = 361$.

The RRBPM is used with the parameters of table VI to compute the smallest 70 energies of CH_3CN. More energy levels could be obtained by using a larger block (B) or by combining the RRBPM with the preconditioned inexact spectral transform method.[78,79] Increasing B would obviously increase the memory cost of the calculation, but with $B = 70$ the memory cost is less than 1 GB. In the direct product basis, storing only one vector would take 1113 GB. We have not incorporated symmetry adaptation and degenerate levels are calculated together. We use the same n_j as in Ref. 50 (but no pruning). The n_j take into account the harmonic frequencies ω_j and some strong coupling terms.

TABLE VI. Parameters for the CH_3CN calculations.

D	12
$n_j, j = \{1, 3, 4, 5, 6, 9, 10\}$	9
$n_j, j = \{2, 7, 8\}$	7
$n_j, j = \{11, 12\}$	27
Reduction rank r	20
N_{ALS}	10
Block size B	70
Maximum N_{pow}	6000
E_{max} estimate	311419
Energy shift σ	170000
N_{diag}	20

The lowest energy levels are given in table VII where they are compared to previous theoretical results. We have determined vibrational assignments in two ways. Note that in the previous section each level was assigned the label of an exact wavefunction but that in this section we assign, in the more usual sense, each level to the $\prod_{j=1}^{D}\theta^j_{i_j}(q_j)$ basis function that most nearly approximates it. First, we use rank reduction to assign. Columns of $\boldsymbol{\mathcal{F}}\ \mathbf{U} = \boldsymbol{\mathcal{F}}^{\mathrm{new}}$ which approximate columns of $\mathbf{V}$, where $\mathbf{HV} = \mathbf{VE}$, have rank Br. Using the ALS algorithm each is reduced to a rank one vector $\prod_{j=1}^{D} f^{(1,j)}_{i_j}$ that approximates a column of $\mathbf{V}$. A level is assigned to $v_j, v_{j'}, v_{j''}, \cdots$ if the factors $f^{(1,j)}_{v_j}$ $f^{(1,j')}_{v_{j'}}$, $f^{(1,j'')}_{v_{j''}}$ etc are

large and all other factors are small. Second, we assign on the basis of components of columns of $\boldsymbol{\mathcal{F}}^{\rm new} = \boldsymbol{\mathcal{F}}\mathbf{U}$. A column of $\boldsymbol{\mathcal{F}}^{\rm new}$ is a linear combination of vectors of the form of Eq. (2) and hence also of the form of Eq. (2). Using the RRBPM it is possible to compute energy levels without ever actually evaluating the sum in Eq. (2), however, to assign we need elements of $\boldsymbol{\mathcal{F}}^{\rm new}$ and must therefore do the sum. To identify dominant components, we calculate only the elements $F^{\rm new}_{b,i_1\ldots i_D}$ suspected to be large. They are those for which $\sum_{j=1}^{n_j} i_j$ is small. When the two assignment procedures yield different results and when no assignment can be deduced from the first we use the second.

Owing to the fact that we use the same basis and same potential as Ref. 50, differences between our energies and those of Ref. 50 are a measure of the accuracy of the RRBPM. The energies of Ref. 50 are obtained using a method that requires about an order of magnitude more memory, but which does not require a simple force-field type potential. Experimental values are also listed in the last column. The RRBPM energies are close to those of Ref. 50. The lowest RRBPM levels in table VII differ from the numbers in Ref. 50 by a few $\rm cm^{-1}$. The higher levels differ more. The levels with the largest errors are: 1779.88, 1780.66 and 2000.43, 2007.90 and the seven highest eigenvalues of the block. The highest eigenvalues could be improved by using a larger B. It is not surprising that eigenvalues at the edge of the block have larger errors. Other large errors could be reduced by increasing r. Some of the eigenvectors corresponding to eigenvalues near the top of the block are nearly linear combinations of the eigenvectors *of different symmetries* computed with the Smolyak quadrature method of Ref. 50. We have not yet implemented a symmetry-adapted rank reduction and therefore the rank reduction breaks the symmetry. In some cases, when levels are close together, the RRBPM energies and those of Ref. 50 are not close enough to match them unambiguously. When this problem occurs we match levels using their assignments, i.e., using the corresponding eigenvectors. Differences between the energies of ref. 50 and experiment are due to the potential. Differences between the energies of ref. 50 and those obtained with the RRBPM are due to the low-rank approximations and could be reduced by increasing r, B, the maximum $N_{\rm pow}$ and N_{ALS}.

V. CONCLUSION

The use of iterative algorithms has opened the door to routine calculation of (ro-)vibrational spectra of molecules with as many as four or five atoms.[5,8–11,18,93,94] The same ideas can be used, with an adiabatic approximation, for Van der Waals complexes with 6 or fewer inter-molecular coordinates.[16,95] Although iterative methods obviate the need to store the Hamiltonian matrix (or even to calculate its matrix elements), application of these ideas to larger molecules is impeded by the size of the vectors that must be stored in memory. Calculations are only "routine" if a product basis is used. With a product basis, the size of a vector scales as n^D. For a $J = 0$ calculation, $D = 12$ for molecule with 6 atoms; storing a vector with 10^{12} elements requires 8000 GB. One way to deal with this impasse is use a contracted basis. Another is to prune a product basis set. In this article we suggest a third approach, based on exploiting the SOP form of the Hamiltonian. At present it can only be applied to SOP Hamiltonians. That is a limitation, but for many molecules with 6 or more atoms for which one wishes to compute a spectrum either the only available potential energy surfaces (PESs) are in SOP form or the PES can be brought into SOP form without making a significant approximation. MCTDH is usually used with a SOP PES, there are however other options.[96,97]

The principal idea of the RRBPM of this paper is the realization that whereas one needs n^D numbers to represent a general function in a product basis, a function that is a SOP can be represented with far fewer numbers. The simplest example is a product of D factors for which one only requires nD numbers, much less than n^D. If the factors of the terms in a SOP representation of a wavefunction are chosen carefully, it should be possible to represent it with a relatively small number of terms. In this article, we show that it is possible to obtain accurate energy levels for a 6-D model problem with $r = 10$, for a 20-D model problem with $r = 20$ and for CH_3CN (12-D) with 299 coupling terms with $r = 20$. This makes it possible to reduce the memory cost of calculations by many orders of magnitude. For the 20-D problem the memory cost is about 1 GB. The RRBPM uses a shifted block power method to make SOP basis functions. Applying the Hamiltonian to a vector necessarily yields the number of terms in the SOP. If this increase were not checked the memory cost of the method would become large. We restrict the number of terms by using a rank reduction idea.[53] A somewhat similar power method idea has been used in Ref. 98. At each stage of the procedure we exploit the SOP structure, e.g., to orthogonalize, to evaluate matrix-vector products etc.

The ideas introduced in this paper can be refined in several ways. Rather than using the block power method to generate SOP functions one could use a better iterative approach. One option is a Davidson algorithm[99], another is a preconditioned inexact spectral transform (PIST) method.[81,100,101] A PIST version would also make it possible to target high-lying levels. Reducing the number of required iterations would reduce the cost of the calculations. Any iterative method whose basis vectors are close enough to the desired eigenvectors to enable rank reduction will be suitable. A symmetry-adapted rank reduction method will obviate symmetry mixing of very nearly degenerate levels. The ALS reduction algorithm is not the most efficient nor the most robust reduction algorithm in the literature. It could be replaced, for example, with a conjugate gradient-based algorithm[77]. The general approach is promising because

TABLE VII. Transition wavenumbers from the ZPE. From left to right: RRBPM results, symmetry, assignment, results from references 50 and 82, experimental values. The zero point energy is 9837.6293 cm^{-1}. RRBPM energies and the energies of Ref. 50 are matched on the basis of assignments. The symmetry is taken from Avila et al[50]. States in a curly bracket are linear combinations (LC) of states in Ref. 50.

Transitions (cm^{-1})	Sym.	Assign.	Ref. 50	Ref. 82	Exp.
361.18, 361.25	E	ω_{11}	360.991	366	362 (Ref. 88), 365 (Ref. 89)
723.37, 724.38	E	$2\omega_{11}$	723.181	725	717 (Ref. 88 and 90)
724.96	A_1	$2\omega_{11}$	723.827	731	739 (Ref. 90)
900.97	A_1	ω_4	900.662	916	916 (Ref 88), 920 (Ref. 89)
1034.50, 1034.55	E	ω_9	1034.126	1038	1041 (Ref 88), 1042 (Ref. 90)
1087.95	A_2	$3\omega_{11}$	1086.554	1098	1122 (Ref. 90)
1088.58	A_1	$3\omega_{11}$	1086.554	1098	1122 (Ref. 90)
1090.75, 1090.85	E	$3\omega_{11}$	1087.776	1094	1077 (Ref. 90)
1260.89, 1261.12	E	$\omega_4+\omega_{11}$	1259.882	1282	1290 (Ref. 91)
1391.76	A_1	ω_3	1388.973	1400	1390 (Ref 92), 1385 (Ref. 91)
1395.74, 1398.24 }	L. C. of E and A_2	$\omega_9+\omega_{11}$	1394.689 (E)	1401	1410 (Ref. 92), 1409 (Ref. 91)
1396.24 }	states	$\omega_9+\omega_{11}$	1394.907 (A_2)	1398 (A_2)	1402 (Ref. 91)
1401.15	A_1	$\omega_9+\omega_{11}$	1397.687	1398	1402 (Ref. 91)
1452.92, 1458.62	E	$4\omega_{11}$	1451.101		
1456.24, 1460.80 }	L. C. of E and A_1	$4\omega_{11}$	1452.827 (E)		
1464.40 }	states	$4\omega_{11}$	1453.403 (A_1)	1467	1448 (Ref. 91)
1483.52, 1483.54	E	ω_7	1483.229	1478	1453 (Ref. 88), 1450 (Ref. 91)
1621.34, 1623.05	E	$\omega_4+2\omega_{11}$	1620.222	1647.0	
1624.05	A_1	$\omega_4+2\omega_{11}$	1620.767	1645.7	
1753.66, 1755.03	E	$\omega_3+\omega_{11}$	1749.530	1766.4	
1759.62, 1760.16	E	$\omega_9+2\omega_{11}$	1757.133	1767.4	
1765.59	A_1	$\omega_9+2\omega_{11}$	1756.426	1761.6	
1779.88	A_2	$\omega_9+2\omega_{11}$	1756.426	1761.6	
1780.66, 1780.86	E	$\omega_9+2\omega_{11}$	1759.772	1769.3	
1786.17	A_1	$2\omega_4$	1785.207	1833.7	
1823.34, 1830.31	E	$5\omega_{11}$	1816.799		
1823.87, 1828.40	E	$5\omega_{11}$	1820.031		
1827.34	A_2	$5\omega_{11}$	1818.953		
1832.19	A_1	$5\omega_{11}$	1818.952		
1845.57 }	L. C. of E and A_2	$\omega_7+\omega_{11}$	1844.258 (A_2),	1838.2 (A_2),	
1846.85, 1849.44 }	states	$\omega_7+\omega_{11}$	1844.330 (E)	1842.2 (E)	
1848.14	A_1	$\omega_7+\omega_{11}$	1844.690	1838.2	
1932.98, 1933.43	E	$\omega_4+\omega_9$	1931.547	1952.3	
1990.80, 1992.59	E	$\omega_4+3\omega_{11}$	1982.857	2015.1	
2000.43	A_2	$\omega_4+3\omega_{11}$	1981.850	2010.3	
2007.90	A_1	$\omega_4+3\omega_{11}$	1981.849	2010.3	
2058.94	A_1	$2\omega_9$	2057.068	2059.0	2075 (Ref. 88)
2066.43, 2068.26	E	$2\omega_9$	2065.286	2067.0	2082 (Ref. 88)
2116.66, 2121.90 }	L. C. of E and A_1	$\omega_3+2\omega_{11}$	2111.380 (E)	2131.3 (E)	
2136.72 }	states	$\omega_3+2\omega_{11}$	2112.297 (A_1)	2130.0 (A_1)	
2126.04 }		$\omega_9+3\omega_{11}$	2119.327 (E)		
2144.72	L. C. for these six	$\omega_9+3\omega_{11}$	2120.541 (E)		
2146.06	states. Two levels	$\omega_9+3\omega_{11}$	2120.910 (A_2)		
2174.61	are missing.	$\omega_9+3\omega_{11}$	2122.834 (E)		
2150.37		$\omega_9+3\omega_{11}$	2123.301 (A_1)		
2153.59 }		$\omega_9+3\omega_{11}$			
2158.75, 2164.96	E	$2\omega_4+\omega_1 1$	2142.614	2199.4	
2210.70	?	$6\omega_{11}$	2183.635 (E)		

its memory cost is low. It might be possible to use similar ideas with the Floquet formalism to solve the time-dependent Schrödinger equation to study a molecule in a strong external electromagnetic field.[102] In such approaches memory cost is a serious problem because the time-dependent Schrödinger equation is solved in an extended Hilbert space containing functions that depend on a time coordinate.[103,104] We have shown that the method can be used to solve the time-independent Schrödinger equation for a molecule with 6 atoms using less than 1 GB. Clearly, it will be possible to compute spectra for much larger molecules.

ACKNOWLEDGMENTS

We thank Gustavo Avila for his help. Calculations have been executed on computers of the Utinam Institute at the Université de Franche-Comté, supported by the Région de Franche-Comté and Institut des Sciences de l'Univers (INSU) and on computers purchased with a grant for the Canada Foundation for Innovation. This research was funded by the Natural Sciences and Engineering Research Council of Canada. We thank James Brown for doing the 6-D direct-product Lanczos calculation.

[1] C. Lanczos, J. of Research of the Nat. Bur. Of Standards **45**, 255 (1950).
[2] J. K. Cullum and R. A. Willoughby, *Lanczos algorithms for large symmetric eigenvalue computations: Vol. I: Theory* (SIAM Classics in Applied Mathematics, 2002).
[3] D. Neuhauser, J. Chem. Phys. **93**, 2611 (1990).
[4] V. A. Mandelshtam and H. S. Taylor, J. Chem. Phys. **106**, 5085 (1997).
[5] R. Chen and H. Guo, J. Chem. Phys. **111**, 464 (1999).
[6] S.-W. Huang and T. Carrington, Chem. Phys. Lett. **312**, 311 (1999).
[7] W. D. H. Koeppel and L. S. Cederbaum, Adv. Chem. Phys. **57**, 59 (1984).
[8] M. Bramley and T. Carrington, J. Chem. Phys. **99**, 8519 (1993).
[9] M. Bramley, J. Tromp, T. Carrington, and G. Corey, J. Chem. Phys. **100**, 6175 (1994).
[10] C. Leforestier, L. Braly, K. Liu, M. Elrod, and R. Saykally, J. Chem. Phys. **106**, 8257 (1997).
[11] F. L. Quere and C. Leforestier, J. Chem. Phys. **94**, 1118 (1991).
[12] N. P. P. Sarkar and T. Carrington, J. Chem. Phys. **110**, 10269 (1999).
[13] G. M. R. Chen and H. Guo, J. Chem. Phys. **114**, 4763 (2001).
[14] J. C. Tremblay and T. Carrington, J. Chem. Phys. **125**, 094311 (2006).
[15] G. M. R. Chen and H. Guo, Chem. Phys. Lett **320**, 567 (2000).
[16] R. Dawes, X.-G. Wang, A. W. Jasper, and T. Carrington, J. Chem. Phys. **133**, 134304 (2010).
[17] G. Avila and T. Carrington, J. Chem. Phys. **135**, 064101 (2011).
[18] J. C. Light and T. Carrington, Adv. Chem. Phys. **114**, 263 (2000).
[19] S. Carter, J. M. Bowman, and N. C. Handy, Theor. Chim. Acta **100**, 191 (1998).
[20] D. M. Benoit, J. Chem. Phys. **120**, 562 (2004).
[21] P. Meier, M. Neff, and G. Rauhut, J. Chem. Theory Comput. **7**, 148 (2011).
[22] R. Dawes and T. Carrington, J. Chem. Phys. **122**, 134101 (2005).
[23] R. Dawes and T. Carrington, J. Chem. Phys. **124**, 054102 (2006).
[24] S. Carter and N. C. Handy, Comput. Phys. Rep. **5**, 115 (1986).
[25] D. T. Colbert and W. H. Miller, J. Chem. Phys. **96**, 1982 (1992).
[26] H.-G. Yu, J. Chem. Phys. **117**, 2030 (2002).
[27] C. Iung, C. Leforestier, and R. E. Wyatt, J. Chem. Phys. **98**, 6722 (1993).
[28] Z. Bacic and J. C. Light, Annu. Rev. Phys. Chem. **40**, 469 (1989).
[29] J. R. Henderson and J. Tennyson, Chem. Phys. Lett. **173**, 133 (1990).
[30] M. Mladenovic, Spectrochim. Acta A **58**, 795 (2002).
[31] D. Luckhaus, J. Chem. Phys. **113**, 1329 (2000).
[32] J. M. Bowman and B. Gazdy, J. Chem. Phys. **94**, 454 (1991).
[33] S. Carter and N. C. Handy, Comput. Phys. Commun. **51**, 49 (1988).
[34] M. J. Bramley and N. C. Handy, J. Chem. Phys. **98**, 1378 (1993).
[35] S. Carter and N. C. Handy, Mol. Phys. **100**, 681 (2002).
[36] H.-D. Meyer, F. Gatti, and G. A. Worth, eds., *Multidimensional Quantum Dynamics: MCTDH Theory and Applications* (Wiley-VCH, Weinheim, 2009).
[37] M. H. Beck, A. Jaeckle, G. A. Worth, and H.-D. Meyer, Phys. Rep. **324**, 1 (2000).
[38] M. J. Davis and E. J. Heller, J. Chem. Phys. **71**, 3383 (1979).
[39] B. Poirier, J. Theor. Comput. Chem. **2**, 65 (2003).
[40] T. Halverson and B. Poirier, J. Chem. Phys. **137**, 224101 (2012).
[41] A. Shimshovitz and D. J. Tannor, Phys. Rev. Lett. **109**, 070402 (2012).
[42] X.-G. Wang and T. Carrington, J. Phys. Chem. A **105**, 2575 (2001).
[43] G. Avila and T. Carrington, J. Chem. Phys. **131**, 174103 (2009).
[44] G. Avila and T. Carrington, J. Chem. Phys. **137**, 174108 (2012).
[45] D. Lauvergnat and A. Nauts, Spectrochim. Acta, Part A , 18 (2014).
[46] M. Bramley and T. Carrington, J. Chem. Phys. **101**, 8494 (1994).
[47] X.-G. Wang and T. Carrington, J. Chem. Phys. **117**, 6923 (2002).
[48] R. A. Friesner, J. A. Bentley, M. Menou, and C. Leforestier, J. Chem. Phys. **99**, 324 (1993).
[49] A. Viel and C. Leforestier, J. Chem. Phys. **112**, 1212 (2000).
[50] G. Avila and T. Carrington, The Journal of Chemical Physics **134**, 054126 (2011).
[51] T. Zhang and G. H. Golub, SIAM J. Matrix Anal. Appl. **23**, 534 (2001).
[52] G. Beylkin and M. J. Mohlenkamp, PNAS **99**, 10246 (2002).
[53] G. Beylkin and M. J. Mohlenkamp, SIAM J. Sci. Comput. **26**, 2133 (2005).
[54] U. Benedikt, A. A. Auer, M. Espig, and W. Hackbusch, J. Chem. Phys. **134**, 054118 (2011).
[55] F. A. Bischoff and E. F. Valeev, J. Chem. Phys. **134**, 104104 (2011).
[56] F. A. Bischoff, R. J. Harrison, and E. F. Valeev, J. Chem. Phys. **137**, 104103 (2012).
[57] E. G. Hohenstein, R. M. Parrish, and T. J. Martinez, J. Chem. Phys. **137**, 044103 (2012).
[58] R. M. Parrish, E. G. Hohenstein, T. J. Martinez, and C. D. Sherrill, J. Chem. Phys. **137**, 224106 (2012).
[59] E. G. Hohenstein, R. M. Parrish, C. D. Sherrill, and T. J. Martinez, J. Chem. Phys. **137**, 221101 (2012).
[60] T. G. Kolda and B. W. Bader, SIAM review **51**, 455 (2009).
[61] F. L. Hitchcock, Journal of Mathematics and Physics **6**, 164 (1927).
[62] L. R. Tucker, Psychometrika **31**, 279 (1966).
[63] W. Hackbusch and S. Kühn, J. Fourier Anal. Appl. **15**, 706 (2009).
[64] L. Grasedyck, SIAM J. Matrix Anal. Appl. **31**, 2029 (2010).
[65] D. Kressner and C. Tobler, Comp. Methods in App. Math. **11**, 363 (2011).
[66] I. V. Oseledets, SIAM J. Sci. Comput. **33**, 2295 (2011).
[67] H. Wang and M. Thoss, J. Chem. Phys. **119**, 1289 (2003).
[68] D. Pelaez and H.-D. Meyer, J. Chem. Phys. **138**, 014108 (2013).
[69] S. Manzhos and T. Carrington, J. Chem. Phys. **125**, 194105 (2006).
[70] S. Manzhos and T. Carrington, J. Chem. Phys. **127**, 014103 (2007).
[71] S. Manzhos and T. Carrington, J. Chem. Phys. **129**, 224104 (2008).
[72] J. C. M. S. E. Pradhan, J.-L. Carreon-Macedo and A. Brown, J. Phys. Chem. A **117**, 6925 (2013).
[73] G. Strang, *Introduction to applied mathematics* (Wellesley Cambridge Press, Wellesley, Massachusetts, 1986).
[74] Y. Saad, *Numerical Methods for Large Eigenvalue Problems*, 2nd ed. (SIAM Classics in Applied Mathematics, 2011).
[75] B. N. Parlett, *The Symmetric Eigenvalue Problem* (Prentice Hall, Englewood Cliffs, NJ, 1980) Chap. 4, (Republished by SIAM, Philadelphia, 1998.).
[76] S. R. Chinnamsetty, M. Espig, B. N. Khoromskij, and W. Hack-

busch, J. Chem. Phys. **127**, 084110 (2007).
[77] M. Espig, W. Hackbusch, T. Rohwedder, and R. Schneider, Numerische Mathematik **122**, 469 (2012).
[78] S.-W. Huang and J. Tucker Carrington, J. Chem. Phys. **112**, 8765 (2000).
[79] B. Poirier and J. Tucker Carrington, J. Chem. Phys. **114**, 9254 (2001).
[80] B. Poirier and J. Tucker Carrington, J. Chem. Phys. **116**, 1215 (2002).
[81] W. Bian and B. Poirier, J. Theor. Comput. Chem. **2**, 583 (2003).
[82] D. Begue, P. Carbonnière, and C. Pouchan, J. Phys. Chem. A **109**, 4611 (2005).
[83] J. K. G. Watson, Mol. Phys. **15**, 479 (1968).
[84] L. Henry and G. Amat, J. Mol. Spectrosc. **5**, 319 (1960).
[85] L. Henry and G. Amat, J. Mol. Spectrosc. **15**, 168 (1965).
[86] H. H. Nielsen, Rev. Mod. Phys. **23**, 90 (1951).
[87] B. Poirier, Private communication (november 2013).
[88] I. Nagawa and T. Shimanouchi, Spectrochim. Acta **18**, 513 (1962).
[89] M. Koivusaari, V. M. Horneman, and R. Anttila, J. Mol. Spectrosc. **152**, 377 (1992).
[90] A. Tolonen, M. Koivusaari, R. Paso, J. Schroderus, S. Alanko, and R. Anttila, Journal of Molecular Spectroscopy **160**, 554 (1993).
[91] R. Paso, R. Anttila, and M. Koivusaari, Journal of Molecular Spectroscopy **165**, 470 (1994).
[92] J. Duncan, D. McKean, F. Tullini, G. Nivellini, and J. P. Pea, Journal of Molecular Spectroscopy **69**, 123 (1978).
[93] B. T. S. Edit Matyus, Gabor Czako and A. G. Csaszar, J. Chem. Phys. **127**, 084102 (2007).
[94] H.-G. Yu and J. T. Muckerman, J. Mol. Spectrosc. **214**, 11 (2002).
[95] X.-G. Wang, T. Carrington, J. Tang, and A. R. W. McKellar, J. Chem. Phys. **123**, 034301 (2005).
[96] U. Manthe, J. Chem. Phys. **105**, 6989 (1996).
[97] F. Huarte-Larraaga and U. Manthe, J. Chem. Phys. **113**, 5151 (2000).
[98] H. Nakatsuji, Accounts of Chemical Research **45** (2012).
[99] E. Davidson, J. Comp. Phys. **17**, 87 (1975).
[100] S.-W. and T. Carrington, J. Chem. Phys. **112**, 8765 (2002).
[101] B. Poirier and T. Carrington, J. Chem. Phys. **116**, 1215 (2002).
[102] S.-I. Chu and D. A. Telnov, Phys. Rep. **390**, 1 (2003).
[103] A. Leclerc, S. Guérin, G. Jolicard, and J. P. Killingbeck, Phys. Rev. A **83**, 032113 (2011).
[104] A. Leclerc, G. Jolicard, D. Viennot, and J. P. Killingbeck, The Journal of Chemical Physics **136**, 014106 (2012).